

# Persistence Mapping Using EUV Solar Imager Data

B. J. Thompson

NASA Goddard Space Flight Center, Code 671, Greenbelt, MD 20771, USA

`barbara.j.thompson@nasa.gov`

and

C. A. Young

NASA Goddard Space Flight Center, Code 670, Greenbelt, MD 20771, USA

Received 15 July 2015; accepted 5 April 2016

## ABSTRACT

We describe a simple image processing technique that is useful for the visualization and depiction of gradually evolving or intermittent structures in solar physics extreme-ultraviolet imagery. The technique is an application of image segmentation, which we call “Persistence Mapping,” to isolate extreme values in a data set, and is particularly useful for the problem of capturing phenomena that are evolving in both space and time. While integration or “time lapse” imaging uses the full sample (of size  $N$ ), Persistence Mapping rejects  $(N-1)/N$  of the data set and identifies the most relevant  $1/N$  values using the following rule: if a pixel reaches an extreme value, it retains that value until that value is exceeded. The simplest examples isolate minima and maxima, but any quantile or statistic can be used. This paper demonstrates how the technique has been used to extract the dynamics in long-term evolution of comet tails, erupting material, and EUV dimming regions.

*Subject headings:*

## 1. Introduction

Many important solar phenomena have dynamic natures that make them difficult to identify and capture. In particular, features that evolve in both space and time require processing techniques that can extract the key physical attributes from increasingly large data sets. We describe a processing technique that is simple to implement, yet captures several important aspects of spatial/temporal evolution. This technique, called “Persistence Mapping,” can be used on its own, or it can provide important pre-processing information for more sophisticated algorithms.

For a data set consisting of  $N$  images with emission values  $I(x, y, t)$ , the Persistence Map  $P_n$  is a function of emission, location and time:

$$P_n(x, y, t_n) = Q(I(x, y, t \leq t_n))$$

where  $x$ ,  $y$  and  $t$  are the spatial and temporal coordinates,  $t_0 < t_n \leq t_N$  are the image sampling times, and  $Q$  is the selection function. The most common (and simplest) forms of  $Q$  are minimum and maximum values, but any quantile or statistic can be used. For a maximum value Persistence Map (such as the ones shown in Figures 3 and 4),  $P_n(x, y, t_n)$  represents the maximum value of the pixel  $(x, y)$  evaluated for the time range  $t_0$  (first image) to  $t_n$  (current image). Similarly, minimum value Persistence Maps are shown in Figures 8, 11, 12, and 13.

Persistence Mapping is an example of a form of information extraction called “image segmentation” (Jain 1989). In general, image segmentation is a process whereby a set of images are partitioned into multiple segments (or sets of pixels, also known as superpixels) to highlight various features. The original pixels are processed using some characteristic or computed property, such as hue, intensity, or texture. The goal of segmentation is to simplify and/or change the representation of an image into something that is more

meaningful and easier to analyze.

Image segmentation algorithms can be grouped into 2 general categories: pixel discontinuities and pixel similarities. Using discontinuity, we partition an image based on abrupt changes, such as edges in an image. This allows the detection of objects and boundaries (lines, curves, edges, etc.) in images, focusing on partitioning an image into regions that are similar according to a set of predefined criteria. Thresholding, region growing, and region splitting and merging are examples of methods in this category (Gonzalez & Woods 2008). This category includes Persistence Mapping. Segmentation algorithms are a useful step in data reduction, but we stress that this method or any image processing method can introduce artifacts that can be misleading, particularly if the user does not examine the original data. Image processing algorithms are not a replacement for, only a complement to, the source data.

While Persistence Maps can capture the evolution of a feature over time, the technique is different from averaging or “time-lapse” integration in that it chooses a single data value to represent location  $(x,y)$  for all times up to  $t_n$  (see Figure 5). If the desired feature or phenomenon is rare or intermittent, averaging or integrating a large data set can significantly decrease the signal relative to the background. The distinction becomes important when the number of images  $N$  becomes large compared to the number of times the feature is observed at a given location; if the feature is only present at location  $(x,y)$  for one image, Persistence Mapping discards the information from the other  $N-1$  images.

There is some ambiguity in the nomenclature, in that our method “applies” persistence to a feature, artificially extending its lifespan, rather than explicitly computing the lifetime (or persistence) of a feature in a dataset. A persistence function was introduced by Fredkin et al. (1985) as a means of characterizing noise in ion-channel activity. Their persistence function is the probability that the current will be at a certain level at a certain point in

time, based on the past variations in the current. The examples we show in this paper are much more simple, and are adapted more for imager data. Additionally, Edelsbrunner et al. (2000) introduce a means of filtering changes in a growing complex and determining the topological persistence of a feature in the presence of growth. Observational maps of the persistence *of* certain phenomena, such as snow cover (Macander et al. 2015), segment the data as a means of determining the longevity of a given phenomenon. Our method does not determine the lifetime of a feature, it does the converse: by “imposing” persistence on an extreme value, it allows the user to identify the growth/expansion of the feature into previously unaffected regions of space. It is an extension of the technique developed by Vecchio et al. (2009), who created minimum and maximum maps of Ca II intensity images to distinguish regions that are influenced by bright network, fibrils, and areas that are influenced by acoustic shocks. Although a feature may not be present in every image, the extrema maps allowed the authors to segment the regions based on their observed behavior.

Of course, the degree of success yielded by the Persistence Mapping technique depends on the ability to optimize the mapping function  $Q$  for a particular feature or phenomenon; not all solar features consistently exhibit clearly identifiable extremes such as minima and maxima. In the next section we present several examples of solar phenomena that fit this criterion, and demonstrate how Persistence Mapping can rapidly distill key characteristics from large data sets.

## 2. Implementation Examples

### 2.1. Maximum value Persistence Mapping of a slowly evolving feature: Comet Lovejoy

Our first example uses data from the Atmospheric Imaging Assembly (Lemen et al. 2012; Boerner et al. 2012) on the Solar Dynamics Observatory (Pesnell et al. 2012). Comet Lovejoy transited the solar corona 2011 December 15-16, and was well observed by investigations on several observatories. Of particular interest is the physics behind the formation of the EUV tail (c.f., Bryans & Pesnell 2012; Schrijver et al. 2012; McCauley et al. 2013; Downs et al. 2013; Raymond et al. 2014).

Comet Lovejoy is a powerful demonstration of the persistence technique because the time history of the comet’s interaction with the corona is key to understanding its behavior. Schrijver et al. (2012), McCauley et al. (2013), Downs et al. (2013) and Raymond et al. (2014) demonstrated how the behavior of the tail presents a unique opportunity to diagnose the magnetic field of the local corona. To study long-term behavior of the tail, Raymond et al. (2014) averaged multiple images during the egress of Lovejoy through the corona; we will demonstrate how Persistence Mapping can be applied instead.

Figures 1 and 2 show the late stages of the inner coronal transit (egress) of Comet Lovejoy. The tail shows variation in structure, with an apparent “kink” shape developing around 00:43:48 UT on 2011 December 16 (the kink is most apparent in the second and third panels of Figure 2). Figure 3 shows the sequence processed with the persistence technique, where  $Q$  is the maximum value evaluated from  $t_0 = 2011\text{-Dec-16}$  at 00:40 UT to the time of each particular image (the top panel of Figure 3 is the original image shown for comparison).

Figure 4 shows the Persistence Maps corresponding to Figure 2, elucidating the tail

emission and tracing out elongated features. The emission from the tail originates near the comet’s egress path, and then spreads in either a northwestward or southwestward prevailing direction; Downs et al. (2013) combined these observations with a magnetic field model to demonstrate how the ionized tail material traces out the local magnetic field lines. The tail consists of individual cores of emitting plasma, originating near the point where the comet nucleus intersects the corona, but then spreading along the connected field lines after ionization. The persistence technique allows the user to clearly identify the flow direction of the comet’s tail emission; the orientation of the mapped features in the top panels of Figure 4 run SE-NW, while the later features run SW-NE. Note that the “kink” in the tail, that is most evident in the second and third panels of Figure 2, corresponds to the change in direction of the striations of emitting plasma in the corresponding panels in Figure 4. The apparent kink in the tail is associated with a change in direction of local magnetic field.

It is important to distinguish persistence methods from “integration” or “averaging,” which are other common ways of combining information from a large number of images. The bottom panel in Figure 3 shows the average value per pixel over the sample period. However, an averaging kernel can dilute information when the majority of the contributing images contain no signal. The emission is rapidly evolving early in the comet’s egress, meaning that the few images that exhibit the tail emission are averaged with the remaining images which show no tail emission. Later in the transit, where the emission lasts for several minutes, the tail is somewhat visible, but not nearly as visible as in the Persistence Map in the center panel. The persistence technique isolates a single value and excludes the others, while averaging produces a less optimal result. Of course, each situation requires a careful consideration of which technique will work best; there are cases where averaging and integration will produce a more ideal result than persistence, such as a collection of individual images with low signal. Persistence may tend to pick out the noise in low-signal images, making it a poor choice of processing method.

Figure 5 illustrates the utility of Persistence Mapping in comparison to multi-image integration. The top frame of Figure 5 is a figure from Raymond et al. (2014), which was produced by summing 26 consecutive AIA 171Å images from 00:44:24-00:49:36 UT. The lower frame of Figure 5 shows a Persistence Map constructed from the same set of images in the top panel. The contrast is dramatically improved, which allows a clearer measurement of the speed, curvature, and variation in emission along and between the striations. The curvature is similar to the arches reported by Schrijver et al. (2012), who explained the phenomenon as variable releases of ionized material with inertia that initially outweighs the Lorentz force of the local magnetic field. Instead of the material passively tracing out a quasi-static field, the magnetic field is distorted by the new mass.

Raymond et al. (2014) analyzed the speed and rate of spreading in a striation (indicated by the black line in the figure) as a means of assessing the role of pickup ions in determining the relative amounts of kinetic and wave energy. The “striations,” and the contrast between them, implied “variations in the Alfvén modes speed of at least factors of two to three.” The measured rate of spreading was consistent with the influence of pickup ions (Zank et al. 2012). In this case, an improvement in the ability to perform these measurements will result in clearer leading edges of the striation motion, and greater accuracy in the contrast between the striations to determine the relative Alfvén speeds.

## **2.2. Minimum value Persistence Mapping of a slowly evolving feature: Coronal dimming**

Our second example involves a coronal dimming feature that evolves slowly with time. Coronal dimmings are known to be good indicators of the site of evacuated material, and possible open field lines, during coronal mass ejections (e.g., Rust & Hildner 1976; Webb et al. 1978; Rust 1983; Sterling & Hudson 1997; Gopalswamy & Hanaoka 1998; Thompson et

al. 2000; Reinard & Biesecker 2008, 2009). They are typically easy to identify, in that they occur in areas of relatively quiet, slowly-evolving Sun and usually last for at least one hour. However, dimmings can extend far from the erupting region, and are sometimes patchy in appearance, with some areas reaching their lowest emission value long before separate but apparently related areas. An important aspect of dimming studies is identifying the full dimming area, and not simply the areas that exhibit a decrease at a given time. We use the dimming/flare/CME event of 2010 November 30 as an example of a “challenging” dimming region.

Figure 6 shows a series of combined wavelength (Red layer = 211Å, Green = 193Å, Blue=171Å) images sampled from 17:05 UT on 2010 November 30 through 2:00 UT 2010 December 1. The first frame, 17:05 UT, is a “pre-event” image, as there is no evidence of the flare and dimming until several minutes later. In Figure 6, we identify several areas that exhibit dimming, including

A: an area that darkens soon after the flare begins, but is later obscured by flare loops

B: an area that darkens soon after the flare begins, but then “recovers” within a few hours

C: a large area opposite Regions A and B, but darkens more gradually

D: a darkened area that appears relatively late in the timeline, after areas A and B have mostly recovered

E: a pre-existing dark region (perhaps coronal hole) that becomes even darker

Given the range of times and locations of the many dimming areas, how can one be certain that all of the dimmings are appropriately captured? The most common method is to perform “base difference” images (Figure 7), where a pre-event image (17:00 UT) is subtracted from each image to highlight only the areas that changed. However, because the different dimming regions grow and recover on different timescales, it is a challenge

to combine images from multiple times to capture the full combined extent of all of the dimming regions. Obtaining the “full” extent is important for those trying to characterize the *total* mass loss (e.g., Harrison & Lyons 2000; Aschwanden et al. 2009), and also for those who wish to understand the complete magnetic footprint of a gradual eruption (e.g., Thompson et al. 2000; Krista & Reinard 2013).

This event is an excellent candidate for the persistence technique because persistence retains the lowest value over the specified timescale, so even if an image starts to “recover” during the sequence it does not impact subsequent images. Figure 8 shows the same images as Figures 6 and 7, processed with the persistence technique where  $Q = \text{minimum value for each individual wavelength from } t_0 < t_n \leq t_N \text{ and } t_0 = 2010 \text{ November 30 at 17:00 UT.}$

There are several things to note about Figure 8 and the associated movie. First, the flare does not appear *at all* in the persistence images and movie, because the technique disregards any pixels that do not decrease in emission. The movie shows dimming regions appearing in the complete absence of any associated flaring activity. (Similarly, a user who is interested in the flare could apply maximum persistence and remove all evidence of dimming.) Second, it is clear from the persistence images that areas A) and B) are in fact one region that evolves inhomogeneously, and that area E) appears to be isolated from the other four. This information can also be derived from examination of the base difference images, but the last panel of Figure 8 illustrates how simply a full dimming map be obtained via Persistence Mapping.

Figure 9 shows the dimming regions as identified in the Persistence Map, with color contours outlining regions of equal intensity decrease quadratically weighted across the three wavelengths (171Å, 193Å, 171Å). The line plot above the figure compares the values derived from pre-event, “base” image subtraction vs. that of the persistence method. The solid black line shows the percentage change in emission summed over the areas marked A,

B, C, D, and E. The dashed line shows the value of the emission in the summed areas as determined by the persistence method. The minimum value of 56% from subtraction can be contrasted with the asymptotic persistence value of 34%; the persistence method estimates 50% more mass loss (44% loss vs. 66%) than the subtraction method.

The reason for the huge difference between the two methods becomes clear when one examines the behavior of the individual regions A - E. The subtraction vs. persistence values for Region A only are shown in red on the diagram. The position on the time axis indicates where the integrated emission in Region A reached its minimum value (18:47 UT). The integrated emission in Region B (blue lines) reaches its minimum value soon after Region A, but it is more than two hours before Region C (purple lines) reaches its minimum. The flare loops have begun to expand into Regions A and B, and these regions have shrunk considerably by then (see Figure 6). Regions D (cyan) and E (green) are delayed even further in reaching their minimum values.

The pie chart superposed on the figure shows the relative contributions of the regions to the total measured dimming given by the black lines. The contributions of Regions D and E are relatively small, but Region C is nearly as large as Regions A and B combined. Because of the significant delay in dimming in Region C relative to Regions A and B, the total mass loss would be underestimated using the integrated subtracted emission (black line), regardless of when it was computed. Therefore, researchers seeking to understand the total mass represented by a dimming, or determine the footprint of all of the magnetic fields involved in dimming, would benefit from using persistence maps instead of traditional subtraction methods.

### 2.3. Minimum value Persistence Mapping of an intermittent feature: Erupting prominence

Our final example details a rapidly evolving phenomenon: falling prominence material viewed in absorption against the bright corona. Figure 10 shows a large prominence on 2011 June 7 that erupts, but a large fraction of the erupting material falls back to the Sun (Reale et al. 2013; Gilbert et al. 2013). The challenge of measuring and analyzing the falling material is complicated by the fact that the overall shape is continually evolving, and individual features are not easily distinguished within the large moving mass. However, there is an advantage in analyzing the motion of these dark moving features in that their absorption properties are not as heavily influenced by thermal evolution as those of bright features. Quasi-stationary bright features can exhibit flows in emission because of changes in thermal properties, whereas the flows of features viewed in absorption are almost entirely due to mass motion. Tracing the locations of dark features is a standard way of determining the trajectory of prominence material (Reale et al. 2013; Gilbert et al. 2013).

In Figure 11 we apply the  $Q$ =minimum technique for the 2011 June 7 eruption as seen by AIA in 193Å, sampling with the highest image cadence (12 seconds) with  $Q$ =minimum and  $t_0$ =2011 June 7 05:00:07 UT. Figure 12 shows a closeup of a portion of the image, highlighting the variation in trajectories. The Persistence Map is able to highlight the various trajectories exhibited by different parts of the prominence. However, the prominence eventually traverses a large fraction of the visible area, and it becomes difficult to distinguish one trajectory from another.

Figure 13 illustrates how the timing and cadence can be chosen to optimize the resultant map. In particular, there is the challenge of determining which trajectories in the SDO images correspond to the trajectories as observed from the STEREO-A EUVI viewpoint, located at .96 AU, 94 degrees from Earth. A 12-second SDO cadence is not ideal

because the location of the prominence does change significantly from one frame to the next. The line is continuous, so individual positions are not easily distinguishable. Additionally, choosing a very long sequence of images is not optimal when multiple trajectories overlap; by fine-tuning  $t_0$  and the sampling time we can isolate individual trajectories more clearly.

In Figure 13 we apply the Q=minimum technique, but instead of sampling with the maximum SDO image cadence (12 seconds), we sample every two minutes up to 11:00 UT, with  $t_0 < t_n \leq t_N$  where  $n = \text{multiples of } 10$  and  $t_0 = 05:00:07$  UT. These maps were used to determine the 3-D trajectories of impacting blobs reported in Gilbert et al. (2013). The temporal evolution of the prominence motion is revealed by the resampling, or “strobing” of images. Although there were many separate trajectories, we projected the locations of a piece of prominence as viewed by SDO to their corresponding lines of sight in the EUVI map. When we required the material in the EUVI map to match the timing and lines of sight of the material in the AIA map, unique trajectories were able to be identified and measured.

Analysis of the unique trajectories in Gilbert et al. (2013) showed that they do not exhibit simple ballistic motion; the features are influenced by the local magnetic field and, similar to the tail of Comet Lovejoy, can be used as an indicator of local magnetic field structure. There are some caveats, of course: there is no expectation that the magnetic field is static, so a blob trajectory is not to be interpreted as a proxy for an extended coronal field line. Additionally, Innes et al. (2012) and Carlyle et al. (2014) have demonstrated that the breakup of the material is consistent with the Rayleigh-Taylor instability, so the flow of material is not exclusively determined by the magnetic field direction. However, the persistence method can highlight sudden changes in prominence direction (see Figure 12) in a field of over 100 other trajectories, allowing the user to isolate the more unexpected types of motion for further analysis.

### 3. Discussion and Conclusions

This technique, while being extremely simple to implement, provides a concise and elegant way to capture the evolution of various solar phenomena. As shown in the examples, the choice of persistence function  $Q$  and sample timing can be tailored to produce a desired result. The persistence technique differs from “long exposure” or “integration” because the latter techniques do not identify and exclude “unwanted” data points. While an integration of 100 images results in 100 values counting equally, the persistence technique completely excludes 99% of the data, retaining only the one value for each pixel that satisfies the  $Q$  criterion. For phenomena evolving in both time and space, Persistence Maps can provide a more distinct and easier way to identify individual features.

Of course, the Persistence Maps become more useful with increasing number of images. This is because the zeroth iteration of the map is simply the original image, while the first iteration differs from the zeroth iteration only in the pixels where the value was exceeded; for pure noise, that would mean 50% of the pixels changed, while 50% retained their original value. The second iteration requires a pixel to exceed two prior values instead of one, resulting in a 33% change on average. It is important to understand the limitations of the maps early in the sequence  $t_0 < t_n \leq t_N$ ; the fraction of pixels that change with each map iteration is on order  $1/n$ , so for low values of  $n$  a large fraction of the pixels will change, while for the 100th image there is only 1% chance of change. This will leave the impression that the features are rapidly evolving early in the series and more static later in the series. Instead, the “settling” effect is (ideally) the result of the algorithm completing the isolation of the desired values. However, if there are only a few images in a dataset, other techniques may more effectively produce the desired result.

Persistence Maps can also be implemented as a “pre-processing step” for more sophisticated analysis; the maps can highlight exactly which subset of the observing area

exhibits a given phenomenon, allowing the user to isolate the regions on which to focus. For the example using coronal dimming images, the maps allow a more complete identification of the regions that exhibit mass loss. For the example using falling prominence material, the technique helps the user identify unique trajectories that can potentially be used as a diagnostic of local magnetic field direction.

As with most methods, artifacts can be misleading. Persistence Maps have the same caveats as most segmentation algorithms: in extracting only a small subset of information from a large dataset, some of the information that is discarded may contain important information, and conclusions based solely on the segmented data may be incorrect. From the map alone, one cannot tell if the feature lasted only one frame or one hundred frames. Therefore, it is recommended that the maps always be viewed in concert with the “normal” images, so the user can view both the development and the decay of a transient phenomenon.

The authors are interested in identifying new applications for the Persistence Mapping technique. We encourage the reader to check <http://sipwork.org/persistence> for updated examples and discussion.

The authors would like to thank Dean Pesnell, John Raymond, Paul Bryans, Michael Chesnes, Mark Cheung, Cooper Downs, Wei Liu, and Leila Mays for discussion and assistance in developing this technique. The authors would also like to acknowledge the referee, Dr. Timothy A. Howard, for his role in improving this paper. This work was supported by NASA Program Element NNH14ZDA001N-GIODDE14.

## REFERENCES

- Aschwanden, M. J., Nitta, N. V., Wuelser, J.-P., et al. 2009, *ApJ*, 706, 376
- Boerner, P., Edwards, C., Lemen, J., et al. 2012, *Sol. Phys.*, 275, 41
- Bryans, P., & Pesnell, W. D. 2012, *ApJ*, 760, 18
- Carlyle, J., Williams, D. R., van Driel-Gesztelyi, L., et al. 2014, *ApJ*, 782, 87
- Downs, C., Linker, J. A., Mikić, Z., et al. 2013, *Science*, 340, 1196
- Edelsbrunner, H., Letscher, D., & Zomorodian, A. 2000, *Proceedings of the 41st Annual Symposium on Foundations of Computer Science*, IEEE Computer Society, 454
- Fredkin, D. R., Rice, J. A., Colquhoun, D., Gibb, A. J. 1995, *Philosophical Transactions: Biological Sciences*, 350, 334, 353
- Gilbert, H. R., Inglis, A. R., Mays, M. L., Ofman, L., Thompson, B. J., Young, C. A. 2013, *ApJ*, 776, L12
- Gonzalez, R. C., & Woods, R. E., 208, “Digital Image Processing”, Prentice-Hall, ISBN # 9780131687288
- Gopalswamy, N., & Hanaoka, Y. 1998, *ApJ*, 498, L179
- Harrison, R. A., & Lyons, M. 2000, *A&A*, 358, 1097
- Innes, D. E., Cameron, R. H., Fletcher, L., Inhester, B., & Solanki, S. K. 2012, *A&A*, 540, L10
- Jain, A. K., 1989, “Fundamentals of Digital Image Processing,” Prentice-Hall, ISBN # 9780133361650
- Krista, L. D., & Reinard, A. 2013, *ApJ*, 762, 91

- Lemen, J. R., Title, A. M., Akin, D. J., et al. 2012, *Sol. Phys.*, 275, 17
- Macander, M. J., Swingley, C. S., Joly, K., & Reynolds, M. K., 2015, *Remote Sensing of the Environment*, 163, 23
- McCauley, P. I., Saar, S. H., Raymond, J. C., Ko, Y.-K., & Saint-Hilaire, P. 2013, *ApJ*, 768, 161
- Pesnell, W. D., Thompson, B. J., & Chamberlin, P. C. 2012, *Sol. Phys.*, 275, 3
- Raymond, J. C., McCauley, P. I., Cranmer, S. R., & Downs, C. 2014, *ApJ*, 788, 152
- Reale, F., Orlando, S., Testa, P., et al. 2013, *Science*, 341, 251
- Reinard, A. A., & Biesecker, D. A. 2008, *ApJ*, 674, 576
- Reinard, A. A., & Biesecker, D. A. 2009, *ApJ*, 705, 914
- Rust, D. M., & Hildner, E. 1976, *Sol. Phys.*, 48, 381
- Rust, D. M. 1983, *Space Sci. Rev.*, 34, 21
- Schrijver, C. J., Brown, J. C., Battams, K., et al. 2012, *Science*, 335, 324
- Sterling, A. C., & Hudson, H. S. 1997, *ApJ*, 491, L55
- Thompson, B. J., Cliver, E. W., Nitta, N., Delannée, C., & Delaboudinière, J.-P. 2000, *Geophys. Res. Lett.*, 27, 1431
- Vecchio, A., Cauzzi, G., & Reardon, K. P. 2009, *A&A*, 494, 269
- Webb, D. F., Nolte, J. T., Solodyna, C. V., & McIntosh, P. S. 1978, *Sol. Phys.*, 58, 389
- Zank, G. P., Jetha, N., Hu, Q., & Hunana, P. 2012, *ApJ*, 756, 21

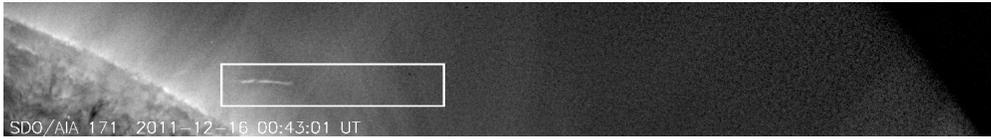


Fig. 1.— Comet Lovejoy as seen in an AIA 171Å image taken 16 December 2011 at 00:43:00 UT. Image has been enhanced radially and with wavelet processing. The northwest limb of the Sun is visible in the lower left corner, and the boundary box indicates the closeup region of Figure 2. The animation of the comet’s passage from 2011 December 16 00:40:00 to 01:00:00 UT is included in the electronic version of this paper.

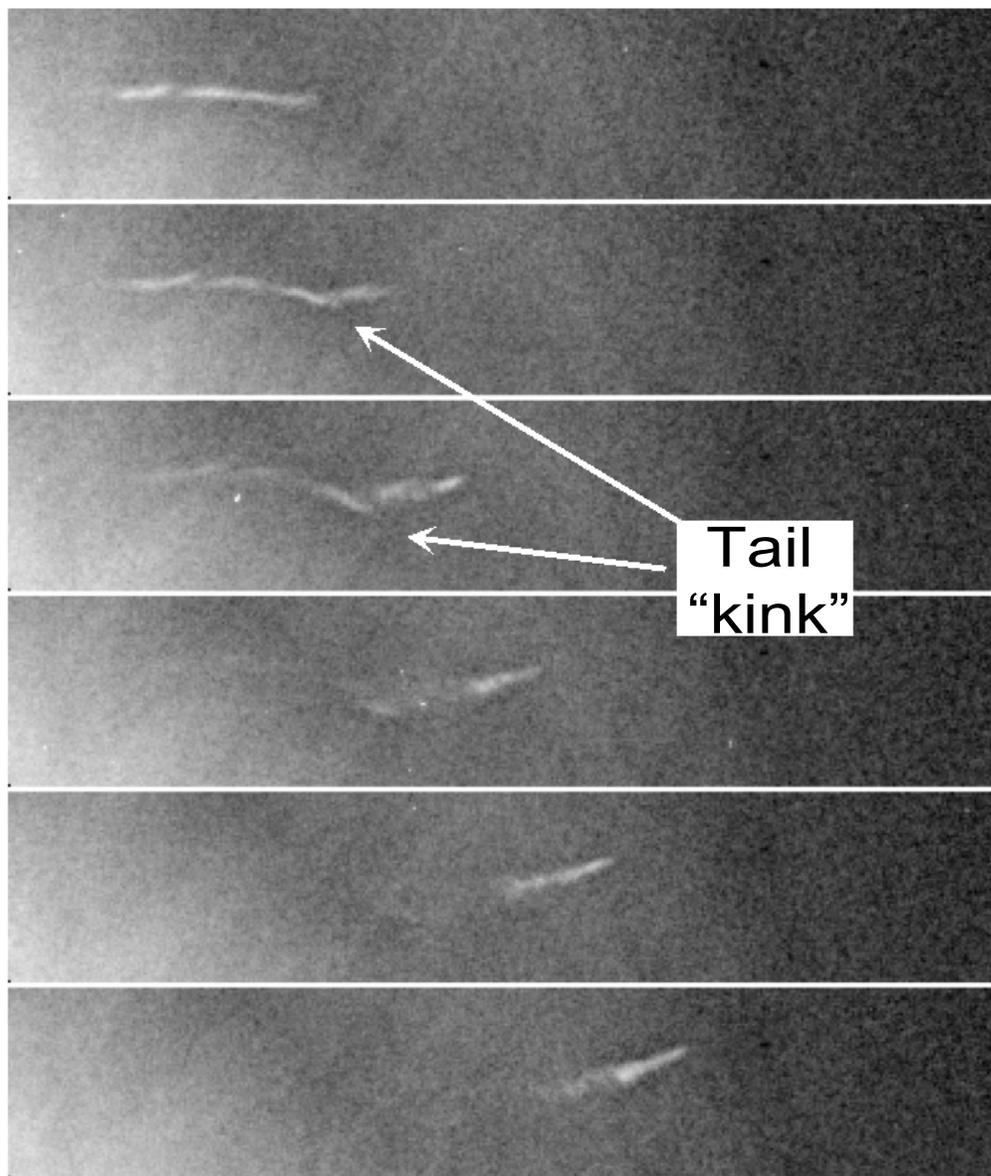


Fig. 2.— AIA 171Å images sampled at 24 seconds (AIA cadence is 12 seconds) from top to bottom: 2011 December 16 at 00:43:00 UT (a closeup of Figure 1), 00:43:24 UT, 00:43:48 UT, 00:44:12 UT, 00:44:36 UT and 00:45:00 UT. Images have been enhanced radially and with wavelet processing.

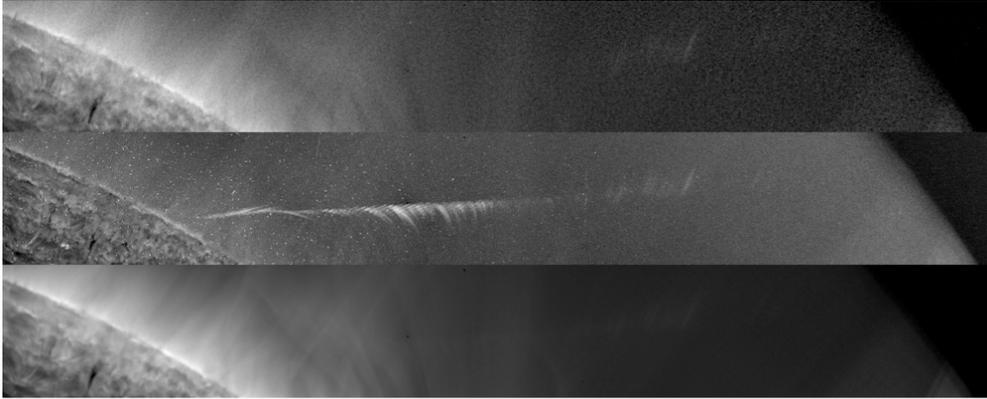


Fig. 3.— Top: AIA 171Å image from 2011 December 16 at 00:57:01 UT. Center: Persistence image at 00:57:00 UT, assembled from 81 AIA 171Å images sampled at a 12-second cadence starting at 00:40:11 UT. Bottom: For comparison, the average pixel value for the entire sequence 00:40 - 01:00 UT is shown. The earliest locations of the tail are not at all visible in the average image, because it is rapidly evolving and the bright features have little impact on the average. The longer-lasting emission later in the egress is faintly visible, but not as clear as in the Persistence Map. All images were enhanced radially and with wavelet processing. The animation of the top and center panels from this figure for the times 2011 December 16 00:40:00 to 01:00:00 UT is included in the electronic version of this paper.



Fig. 4.— The top five frames show the images from Figure 2 (corresponding to times 00:43:00 UT, 00:43:24 UT, 00:43:48 UT, 00:44:12 UT, 00:44:36 UT) processed using the persistence technique, where  $Q$  is the maximum value function for  $t_0 < t_n \leq t_N$  and  $t_0 = 00:40:11$  UT. The top 5 panels were assembled using AIA 171Å images sampled at a 12-second cadence starting at 00:40:11 UT, up to the time corresponding to the panel. The sixth panel shows the persistence values evaluated at 00:49:00 UT to show the later progress of the comet. The lateral striation effect is due to the sampling period of the images.

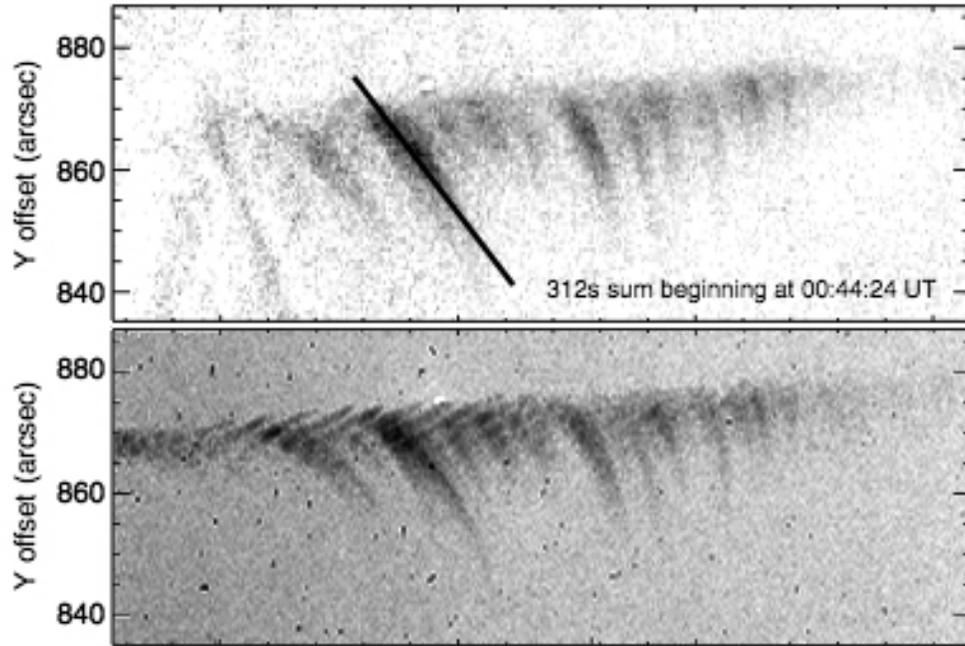


Fig. 5.— The top frame shows a sum of 26 consecutive AIA  $171\text{\AA}$  images of Comet Lovejoy from 00:44:24 - 00:49:36 UT (adapted from Raymond et al. (2014)). The black line indicates a reference direction for a single striation. The bottom frame shows the Persistence Map of the same series of images. The increased contrast from the Persistence Map affords an improved identification of the structure of the striations.

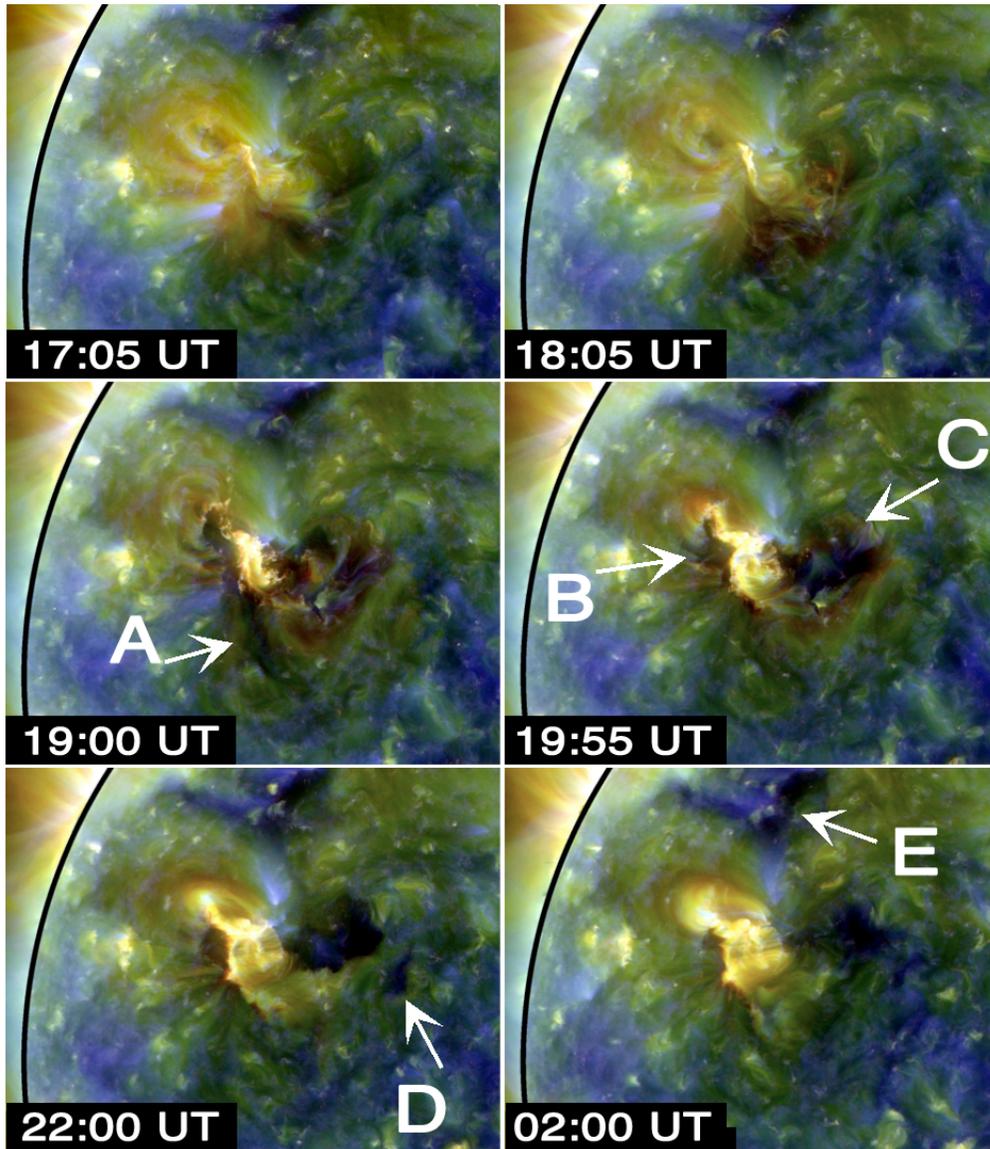


Fig. 6.— AIA combined wavelength images (Red=211Å, Green=193Å, Blue=171Å) for the dimming/flare/CME event on 2010 November 30 17:05 UT, 18:05 UT, 19:00 UT, 19:55 UT, 22:00 UT, (December 1) 02:00 UT. Regions A - E all exhibited dimming at some point during the event, but the much of the early dimming has disappeared by the time region ‘E’ appears. The animation of this figure from 2010 November 30 17:00 - 2010 December 1 02:55 UT is included in the electronic version of this paper.

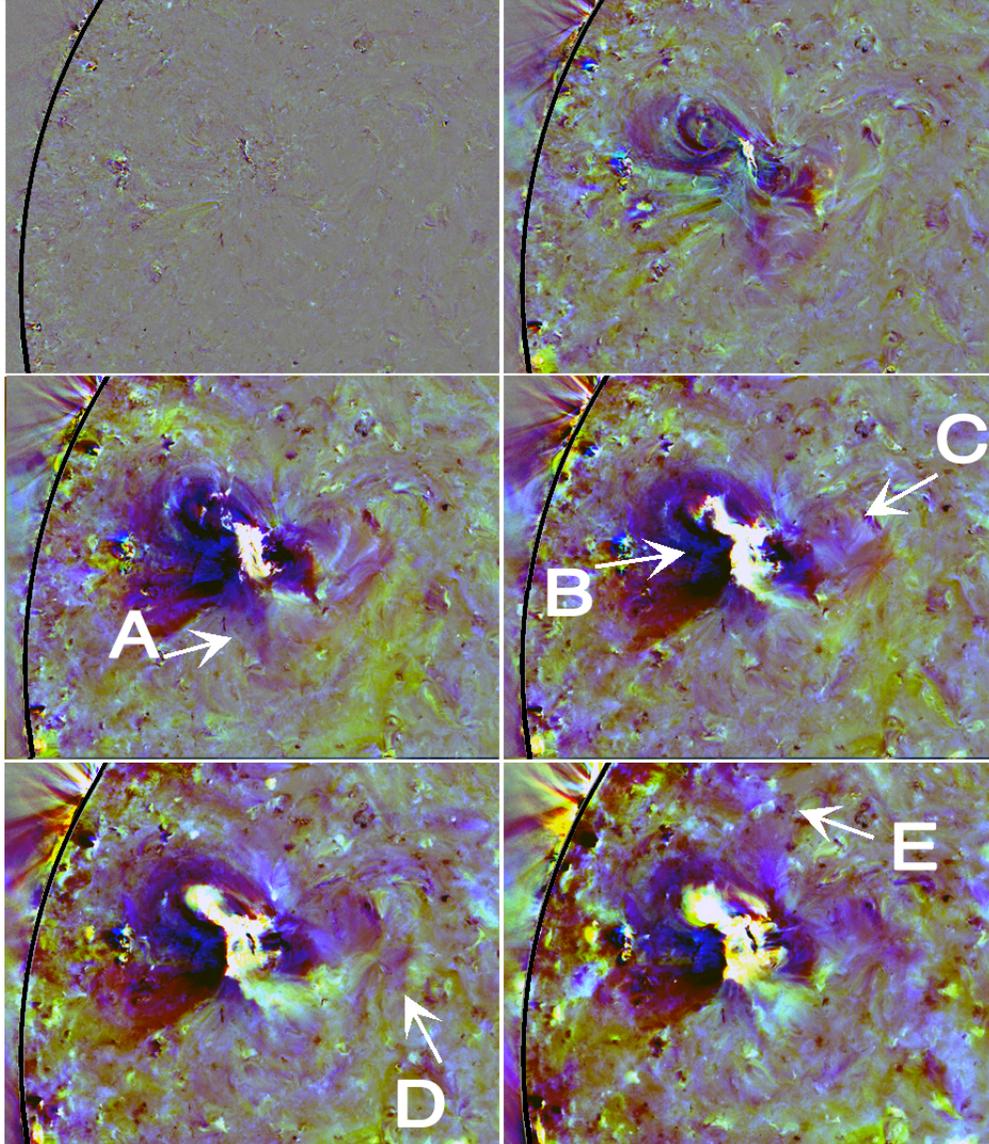


Fig. 7.— AIA combined wavelength images from Figure 5 with “base images” at 17:00 UT subtracted off (i.e. Red layer is 211Å images with the 17:00 211Å image subtracted from it, Green is 193Å images with the 17:00 193Å image subtracted from it, and same for Blue and 171Å). The animation of this figure from 2010 November 30 17:00 - 2010 December 1 02:55 UT is included in the electronic version of this paper.

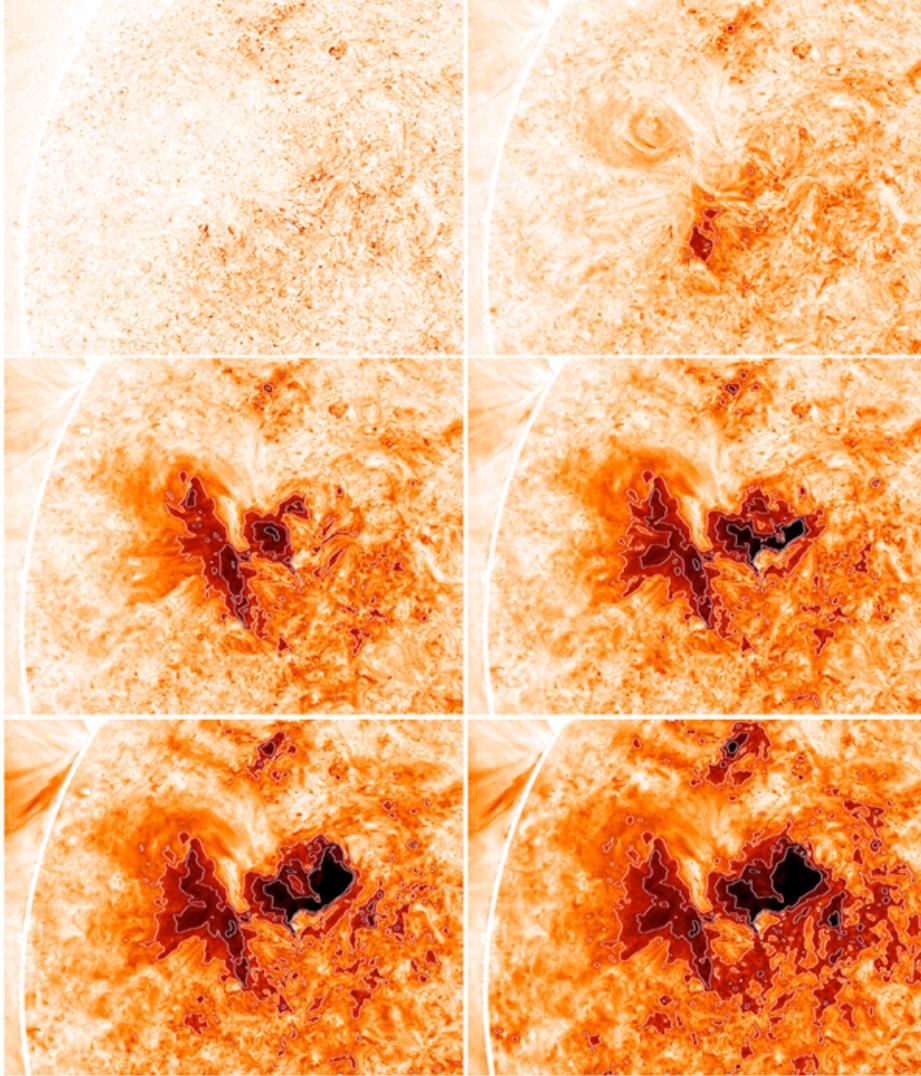


Fig. 8.— AIA combined wavelength images processed with the persistence technique. Each wavelength was processed independently using the persistence technique, with  $Q = \text{minimum value per wavelength from } t_0 < t_n \leq t_N$  and  $t_0 = 2010 \text{ November } 30 \text{ at } 17:00 \text{ UT}$ . Contours outline intensity decrease quadratically weighted across the three wavelengths (171Å, 193Å and 211Å). The animation of this figure from 2010 November 30 17:00 - 2010 December 1 02:55 UT is included in the electronic version of this paper.

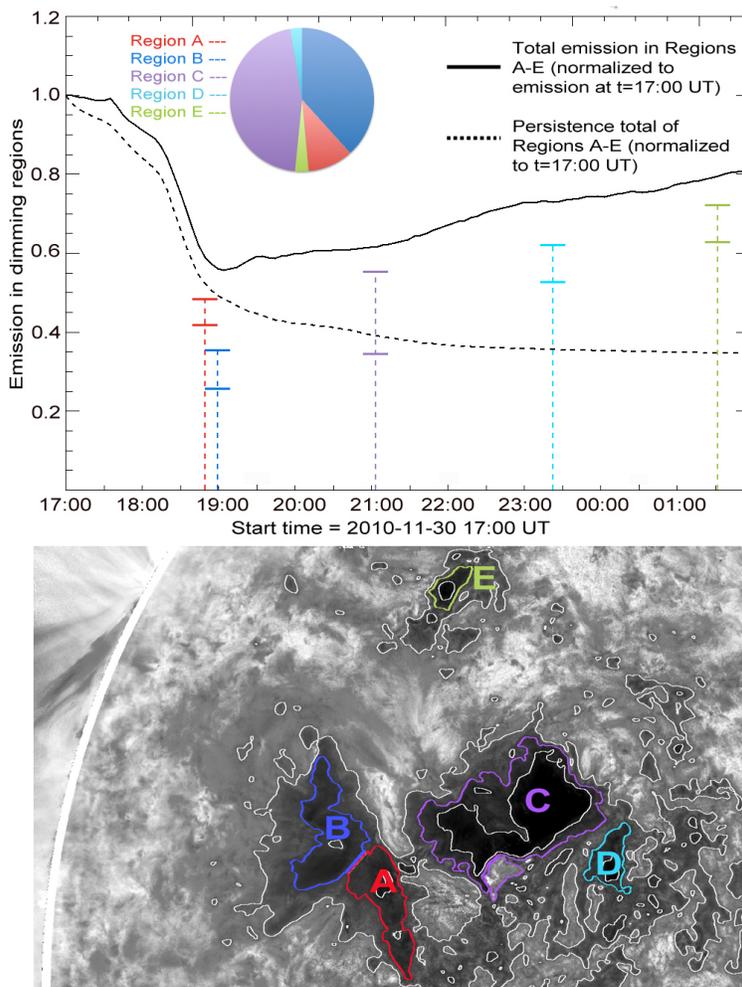


Fig. 9.— The final Persistence Map in Figure 8 is shown in the lower panel, with color outlines corresponding to Regions A - E indicated in Figure 6. The line plot above the figure compares the values derived from pre-event “base” image subtraction vs. that of the persistence method. The solid black line shows the percentage decrease in integrated emission of the combined areas marked A, B, C, D, and E, normalized to the pre-event value. The dashed black line shows value of emission decrease from the same combined regions as determined by the persistence method. The subtraction vs. persistence values for each of the individual regions A - E are shown using their associated color (A = red, B = blue, etc.). The position of the region’s values on the time axis indicates where the integrated emission in the region reached its minimum value. The pie chart superposed on the figure shows the relative contributions of the regions A - E to the total measured dimming.

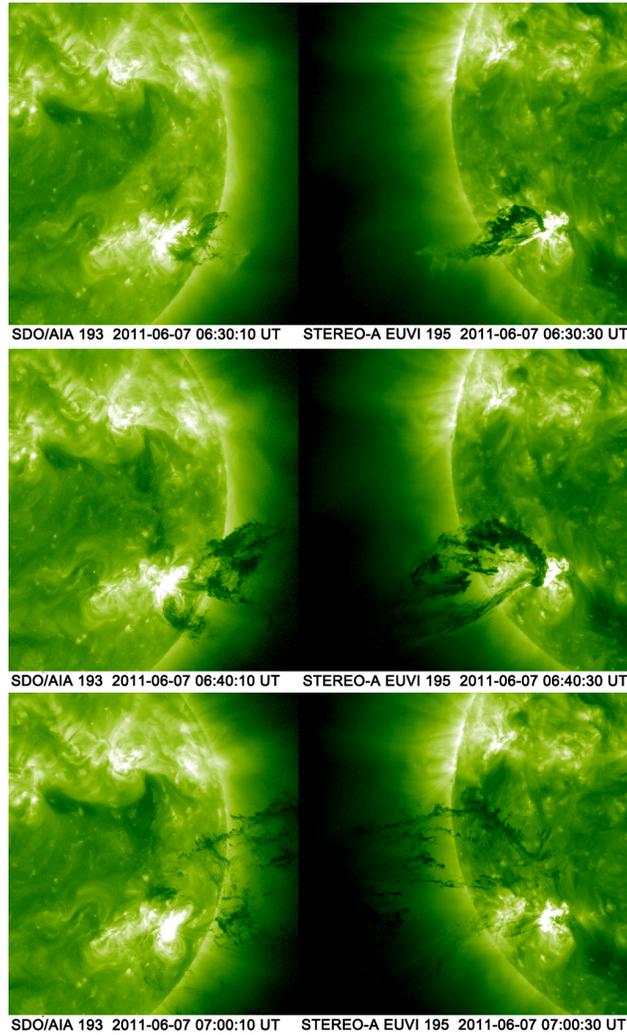


Fig. 10.— Erupting prominence observed from two viewpoints: SDO 193Å (left panels) and STEREO-A EUVI 195Å (right panels). Note that the color table for SDO 193Å has been altered to match the EUVI color table to facilitate comparison. The animation of this figure from 2011 June 7 05:40 - 11:55 UT is included in the electronic version of this paper.

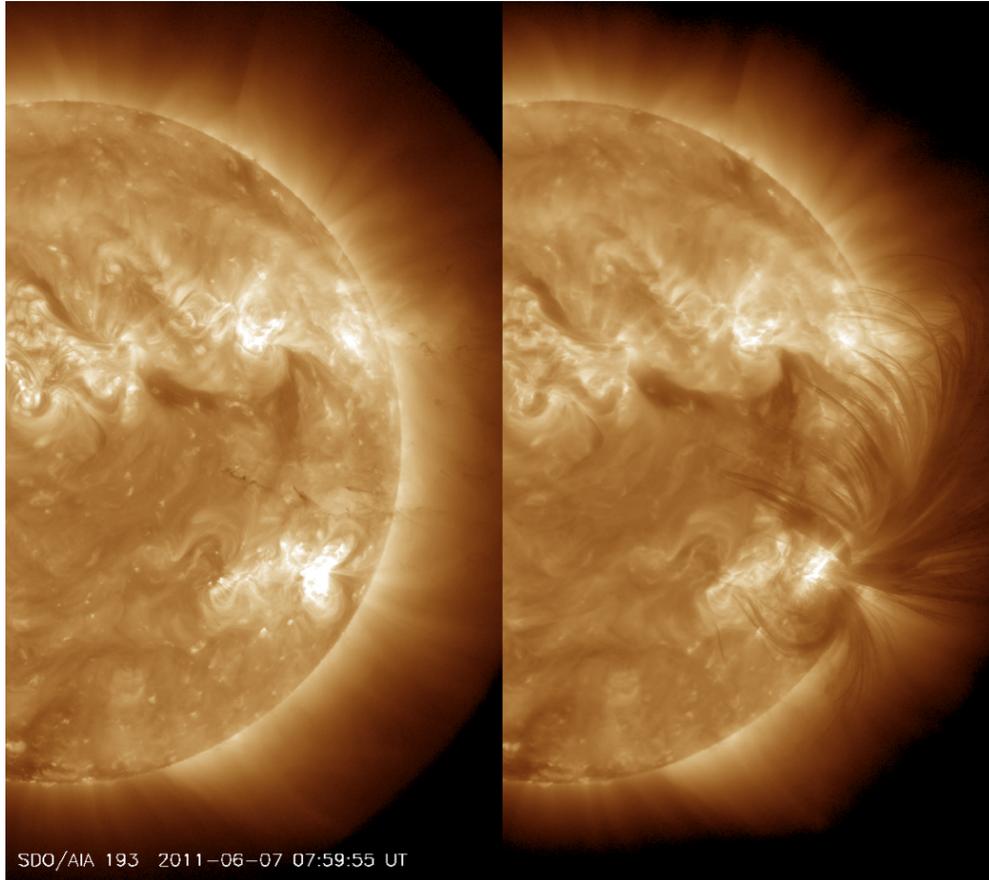


Fig. 11.— Persistence Map evaluated at 2011 June 7 07:59:55 for the SDO 193Å images with  $Q = \text{minimum value per wavelength from } t_0 < t_n \leq t_N \text{ and } t_0 = \text{2011 June 7 at 05:40 UT.}$  The animation of this figure from 2011 June 7 05:40 - 08:00 UT is included in the electronic version of this paper.

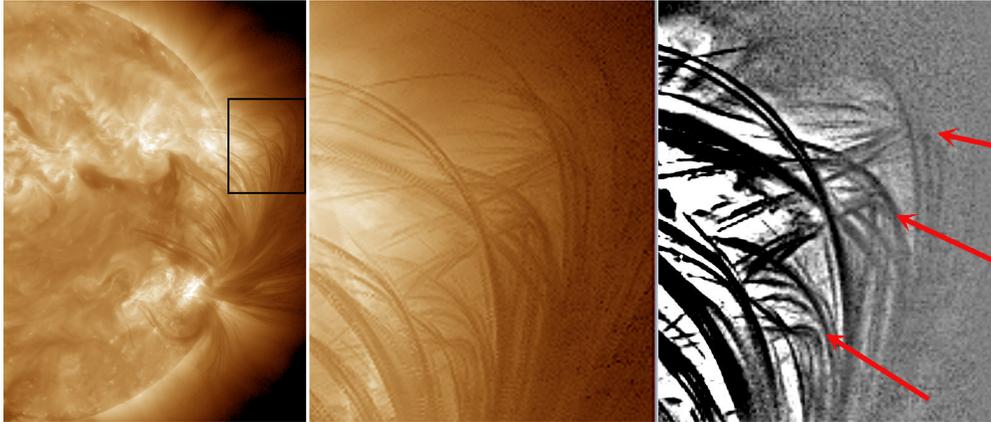


Fig. 12.— A closeup view the trajectories in the Persistence Map shown in Figure 11. The box in the left panel shows the location of the closeup views in the second panel. The third panel is an enhanced version of the second panel. The red arrows indicate sharp bends in the trajectory of some pieces of prominence, which contrast from the majority of the rest.

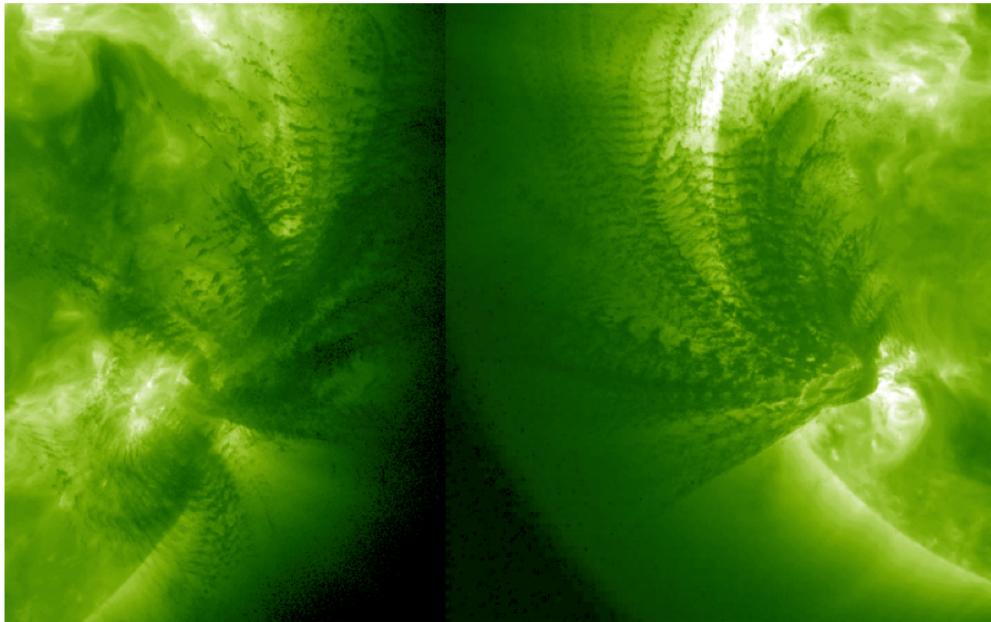


Fig. 13.— Persistence Maps of SDO and STEREO-A EUVI images for the fields of view shown in Figure 10, “strobed” at a cadence of 2 minutes, as opposed to the 12-minute cadence in Figure 11.